



Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

#2
12/21/00
JC825 U.S. PTO
09/676997
09/29/00

Bescheinigung

Certificate

Attestation

Die angehefteten Unterla-
gen stimmen mit der
ursprünglich eingereichten
Fassung der auf dem näch-
sten Blatt bezeichneten
europäischen Patentanmel-
dung überein.

The attached documents
are exact copies of the
European patent application
described on the following
page, as originally filed.

Les documents fixés à
cette attestation sont
conformes à la version
initialement déposée de
la demande de brevet
européen spécifiée à la
page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

99480110.8

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

I.L.C. HATTEN-HECKMAN

DEN HAAG, DEN
THE HAGUE,
LA HAYE, LE

05/01/00

This Page Blank (100)



Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

Blatt 2 der Bescheinigung
Sheet 2 of the certificate
Page 2 de l'attestation

Anmeldung Nr.:
Application no.: 99480110.8
Demande n°:

Anmeldetag:
Date of filing: 21/10/99
Date de dépôt:

Anmelder:
Applicant(s):
Demandeur(s):
International Business Machines Corporation
Armonk, N.Y. 10504
UNITED STATES OF AMERICA

Bezeichnung der Erfindung:
Title of the invention:
Titre de l'invention:

System and method for enabling remote surveillance of ATM network switching node ports

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s) revendiquée(s)

Staat:
State:
Pays:

Tag:
Date:
Date:

Aktenzeichen:
File no.
Numéro de dépôt:

Internationale Patentklassifikation:
International Patent classification:
Classification internationale des brevets:

/

Am Anmeldetag benannte Vertragsstaaten:
Contracting states designated at date of filing: AT/BE/CH/CY/DE/DK/ES/FI/FR/GB/GR/IE/IT/LI/LU/MC/NL/PT/SE
Etats contractants désignés lors du dépôt:

Bemerkungen:
Remarks:
Remarques:

This Page Blank (uspto)

**SYSTEM AND METHOD FOR ENABLING
REMOTE SURVEILLANCE
OF ATM NETWORK SWITCHING NODE PORTS**

5 Field of the Invention

The present invention relates generally to ATM (Asynchronous Transfer Mode) multi-node networks and more specifically to a system and method for enabling the surveillance, through the nodes and switching pieces of equipment forming the
10 network, of any switching node port from any convenient remote access point.

Background of the Invention

Broadband ISDN (Integrated Services Digital Network) was an attempt to set up a single unified, worldwide high-speed network in place of the multiplicity of existing networks for different applications. The new, universal network was intended to be able to take over, on one hand, the functions of current speech, data and television networks and, on the other hand, to provide enough scope for the implementation of future communications technologies. The first work on standards for this universal network of the future was begun by CCITT (International Telegraph and Telephone Consultative Committee) in the late 80's, under the heading of B-ISDN. It is based on ATM which is a data transmission technique belonging to the family of cell switched systems (cell relay). Unlike packet-switched systems, in which data packets of variable length can be multiplexed over a line interface, the length of a cell relay data packet is fixed and simply referred to as a cell. ATM is a specific implementation of cell relay and an integral element of the CCITT specification for B-ISDN. ATM, as the name may suggest, uses an asynchronous time division multiplexing scheme so that the data streams to be transmitted are converted into fixed-size cells and transferred asynchronously over a same physical medium between network nodes. The allocation of the units of information to the different transmission channels is carried out using numerical channel identifiers attached to each cell under the form of a VPi (Virtual Path identifier) and a VCi (Virtual Channel identifier). ATM is a cell switching technology i.e., there are cell-switch units at nodes of the network, in charge of routing cells towards their final destination and which may have to buffer the cells before transmission to a next node. Cell switching implies that an end to end transmission path is set up and must be in existence before any data can be transferred between sender and receiver even though this path is a virtual channel which is only used when required. Thus,

cell-switched networks, can allocate unused transmission capacity to other virtual connections taking advantage of the greatly varying bandwidth requirements in data communications. As cell-switched networks can only create virtual transmission

5 links (over a common physical transmission line), it is possible for cells to be lost if the memory capacity of the switching nodes is exceeded. This transfer procedure, which first requires a (virtual) link to be set up between the users, is said to be connection-oriented (as opposed to connectionless

10 like e.g., IP the Internet Protocol). Links can be PVC (Permanent Virtual Connection) set up once for all or dynamically, on demand, and referred to as SVC (Switched Virtual Connection). It is only after such a connection has been established, be it permanent or dynamic, that sending station

15 can forward its data, under the form of cells, to the receiving station which receives them in the same order. However, if problems occur during the transfer of the data such as a buffer overload, here above already mentioned, resulting in cells discarding, or the receipt of faulty or misdirected

20 cells this is reported immediately to the other end, which can react appropriately generally by repeating the transmission of the cells forming a message so that this can be handled by the error recovery routines of the higher level protocols.

Thus, ATM switching units or just switches are key

25 elements in any B-ISDN network. The fact that all ATM cells are the same size is exploited to implement multiport switching fabrics (typically 16x16 or 32x32) based on various efficient architectures which, when combined with the latest sub-micron fabrication processes, allow to reach very high

30 performances and aggregate throughputs which must be expressed in terabits per sec (10^{12} /sec). Then, the task of the switching fabric and switches in general is, at each node of a network, to provide transmission paths between the input ports and output ports requested at any given time in such a way that

35 the fewest possible conflicts occur. For example, an internal conflict to a switch fabric may occur if two or more cells are

competing for the same output port at the same time. Although, every effort is made during the architecture and design phases of switches and cell-networks in general to put in place all the mechanisms necessary to handle the flows of data without
5 conflicts it is inevitable that congestions may occur occasionally in a node when the overall traffic is increasing or just because a unusual large portion of the traffic is flowing towards the same node at a given instant. Also, flaws and shortcomings that may be present in some of the numerous
10 hardware and software pieces that together implement a network, although not serious enough to prevent the network from operating reasonably well and, in any case, at the satisfaction of the end users, may also trigger occasionally some of the problems already mentioned here above such as cells
15 discarding.

Therefore, it is of utmost importance for those in charge of running and maintaining such an ATM network to be able to watch it and analyze it, if not constantly, at least any time it is necessary, from any convenient access point, without
20 having to disturb user traffic whatsoever. Moreover, because ATM networks are preferably used nowadays to implement only the high performance backbone of many medium or large networks (on contrary of the initial expectation that ATM would become the universal means to transport data, voice and all multime-
25 dia information, up to desktop) the focus is even more highly on performances and quality that only ATM networks can offer by design thus, must be enforced with the proper tools to make sure that the high level of expectation on the quality of ATM network is indeed delivered.

30

Object of the Invention

Thus, it is a broad object of the invention to permit that switch ports, at any node of an ATM network, be observable from any convenient remote access point.

5 It is a further object of the invention to allow that all unexpected and faulty traffic observed at any given entry port of a switching node, irrespective of their actual destination and channel identifier, be collected and forwarded to a predefined remote port for investigation.

10 Further advantages of the present invention will become apparent to the ones skilled in the art upon examination of the drawings and detailed description. It is intended that any additional advantages be incorporated herein.

Summary of the Invention

A method and a system enabling remote surveillance of any entry port to any switching node of an ATM cell-relay network are disclosed. Such an ATM cell-relay network usually
5 comprises several switching nodes each having a plurality of ports. The method first assumes that a path is set up, through the ATM network, between an entry port and a remote observation point. Then, all cells of the incoming traffic, entering entry port to be watched, are duplicated and marked. After
10 which they are transported, following the path, up to the observation point. The preferred method of marking all duplicated cells consists in reserving one bit of the cells channel identifier to unambiguously distinguish them from the regular cells. This bit is asserted in every duplicated cell so as
15 they are kept unaltered and process expedited in the intermediate switching nodes en route to the observation point.

The method and system of the invention permit that any entry port of an ATM network, that may well span over large geographic areas, be conveniently observable and analyzed
20 transparently i.e., without disturbing users traffic, from a remote location so as network can be maintained and run trouble free.

Brief Description of the Drawings

- Figure 1** depicts the key component of any ATM network i.e., the switch.
- Figure 2** is an example of a large ATM network.
- Figure 3** is a closer look to a switch comprising a switch-fabric with adaptive blades interfacing data communications lines.
- Figure 4** illustrates how cells are replicated either within switch fabric or in blades.
- Figure 5** shows structure of an ATM cell and the preferred method of marking them.
- Figure 6** depicts the process of a cell entering a blade.
- Figure 7** depicts the process of a cell exiting the switch fabric to a blade.

Detailed Description of the Preferred Embodiment

Figure 1 illustrates the key component of all ATM networks i.e., the switch. In this figure, as an example, a 16x16 port switch [100] is shown. Each of the 16 ports e.g., port [110], is a bi-directional port capable of receiving and sending ATM cells [120] carrying a connection identifier under the form of a VPi (Virtual Patch identifier) [121] and a VCi (Virtual Channel identifier) [122] part of the header [125] of every cell. Switch performs two basic jobs. First, it identifies the connection identifier of every incoming cell from the VPi and VCi fields here above mentioned. Then, it transports cells from input ports to output ports (e.g. [130]) one step forward along the predetermined path towards their final destinations. The actual transport is carried out via the switching fabric [140] providing dynamic transmission paths between the input ports and the output ports requested at any given time in such a way that the fewest possible conflicts occur like when two cells, from two different input ports, are competing for the same output port at the same time. There are numerous architectures that have been proposed and implemented which all tend to reduce the transmission delay while increasing the aggregate throughput that can be handled through a switch. Nowadays, this latter parameter must be expressed in terabits/sec or 10^{12} bits/second. At the same time switch designs tend to drastically limit the number of cells that must be discarded due to conflicts of the kind mentioned above. Cell loss rate as low as 10^{-10} are reported i.e., less than one cell out 10^{10} cells going through a switch need to be discarded under nominal conditions. Therefore, ATM switches indeed allow to reach the high level of performance required to face the exponential growing demand for bandwidth so as to transport all sorts of traffic such as pure data, voice, still images and video on a unified network.

Figure 2 shows a network built around a high performance ATM backbone comprising in this particular example five switches like [201] or [202]. Switches are interconnected through high speed links such as [210] e.g., optical lines conforming to the well known SDH (Synchronous Digital Hierarchy) or SONET (Synchronous Optical Network) transmission standards respectively defined by the telecommunication standardization sector of the International Telecommunications Union or ITU-T in various G-series standards and by the American National Standards Institute or ANSI e.g., in T1.105. Those two similar standards have transmission rates that are compatible at 155.2 Mbps (referred to as STM-1 and STS-3 for SDH and SONET respectively), 622 Mbps (STM-4 and STS-12) and 2.48 Gbps (STM-16 and STS-48). Thus, allowing to build very high performance backbone networks [220] capable of moving huge amount of data over large distances that may well span over states and countries and allow to interconnect thousands to millions of end-users like the ones on a Token Ring LAN [230], somewhere in the north part of Europe, that may have to get connected with southern users on an Ethernet LAN [240] while both teams may need to access resources from a third remote place e.g., [250]. It is of the utmost importance that such a backbone be operational 24-hour a day 7-day a week. Various techniques, known from the art, based on redundancy, are employed to be as close as possible to a 100% availability. Hence, it is often the case that inter switch transmission lines are actually duplicated so as a hot standby link is always ready to take over a failing one as show in [211]. Therefore, this kind of network must be managed carefully and cannot be left unattended since, if it would collapse, a huge number of users could be impacted simultaneously and critical transactions impaired. This must be conveniently carried out from any access point to the backbone from which the network administrator has decided to watch and manage the entire network. Then, from there, it must be ideally possible to watch a remote switch port such as [231] (or any other entry

port to any of the switching nodes forming the network [220] irrespective of their location at the periphery of the network or within the network) from a distant location [260] where measuring equipment has been installed on a port [261] of switch [201] for the purpose of snooping all the traffic entering through port [231] and generated by the group of users connected on Token Ring [230]. This must be achievable through the direct link [210] between switches [201] and [202] or through any alternate path [215] encompassing other switches of the backbone if more appropriate. This must be possible without necessitating the addition of any hardware, if one excepts the tooling equipment in [260] and, obviously, without requiring the installation of dedicated lines between switches. This latter requirement is in practice, most often, not to say always, automatically fulfilled in such backbone where inter-switch links are very high speed links sometimes duplicated for the reason mentioned above and whose spare bandwidth can be indeed utilized occasionally for the surveillance of the network. Therefore, the invention permits to meet all of the above requirements and is further described in the following figures through a preferred embodiment.

Figure 3 is a closer view on a few ports to a switch fabric [300]. Each active port of a switch is equipped with a port adapter here after referred to as an adaptive blade or just a blade. Three of them are shown namely [310], [320] and [330] respectively connecting to the bi-directional switch fabric ports [301], [302] and [303]. To the other end, blades interface data communications lines such as [317] in a variety of combinations including the number of lines, their speeds, the standards they are complying to and the types of physical interface e.g., optical or electrical. A typical example could be a blade interfacing sixteen optical STM-1/STS-3 communications lines [315] at 155.2 Mbps i.e., compatible with the European SDH and US SONET standards. Such a blade must be able to accommodate the incoming and outgoing traffic of the

sixteen lines to a single switch-fabric port thus, having to handle an average bi-directional traffic up to 2.4 Gbps .

Therefore, the role of a blade such as [310] is to concentrate and dispatch traffic from/to several lines from/to a higher speed switch fabric port so as to expand as much as possible the connectivity of the switch [300] and take full advantage of its intrinsic performances. Hence, in practice, a switching node is comprised of a switching fabric having ports to which blades are connected to allow adaptation to one or more transmission lines which become, in turn, ports of the switching node. Thus, this term (i.e., port) is employed in the following description either to refer to the switching fabric ports themselves or to the ports of the switching node i.e., the end points of the transmission lines connected to the switch blades. This however cannot be ambiguous from the context.

Finally, in turn, a 155 Mbps line [317] could be the uplink to/from e.g., a few local area networks, through any convenient peripheral box as shown in figure 2 [240] further expanding the connectivity of a switching node.

Then, a blade is comprised of a receive side [311] handling all the incoming traffic received from the sixteen lines, and a transmit side [312] handling all the outgoing traffic to be dispatched over the sixteen lines [315] from the switch-fabric port.

The chief function of the receive side [311] is to inspect each received cell thus, enforcing policing and performing the verification of the connections according to the Quality of Service (QoS) defined for each connection possibly tagging or discarding nonconforming cells. The main task being to find out the destination of the incoming cell i.e., the port and blade it should be sent so as cell moves along the path assigned by the network manager (ATM is a connection oriented protocol which assumes that a path exists before traffic can start). This is found from a connections table [316] that keeps track of all active virtual connections supported at any given time by the blade. Therefore, whenever a cell is received,

this table is interrogated so as to determine through which outgoing port it must exit thus, appending to the incoming cell the necessary information in order it can find its way out through the switch fabric itself. Depending on the switch fabric design and implementation numerous solutions are possible to achieve this. Irrespective of any particular implementation incoming cells are eventually inputted to the switch fabric through the port to which blade is connected [301] then switch fabric manages to forward the cell to the transmit side of the appropriate target blade. At this point it is worth noting that more than one switch fabric output port may need to be specified in the information added to the cell because some cells must be forwarded to several ports in case of multicast. This function is generally supported in the switch fabric itself which is capable of replicating the incoming cells in order they are dispatched over several ports when necessary. Moreover, another level of multicasting may have to be performed in the transmit side of the blade itself. For example, in a multi-line blade [320], similar to [310], a cell getting out of the switch fabric port [302] may have to be replicated several times so as it is possibly dispatched over the sixteen lines this blade interfaces or even multiple times over the same line. To illustrate this two cases of duplication are shown. A first one within the switch fabric itself showing that the incoming flow of cells [318] may have e.g., to be duplicated over port [302] and [304]. This flow may need to be duplicated again, in the transmit side of blade [320] e.g., over lines [321] and [324] when multicasting requires that destinations accessible through those lines need to receive the multicast flow of cells. This information is found through the interrogation of a local connection table [326] for each cell received through the switch fabric port [302] in a manner similar to what is done for all incoming cells through line [317] and table [316].

Figure 4 starts describing the preferred embodiment of the invention which is broadly referred to as port snooping in the rest of the description. Performing snooping of a node port i.e., of all the incoming traffic arriving through a line such as [417], so that it can be remotely observed, first consists in marking on-the-fly all the cells entering the blade from this port then, transport a copy of them up to the observation point possibly through many intermediate other nodes. When snooping mode is entered a snoop source point of interest is chosen from which a path must be setup up to the observation point. Like with all regular ATM connections the path is set up in advance and must have a reserved bandwidth compatible with the actual throughput of the snooped port so that its traffic can be safely transported and observed from a remote location. This was previously illustrated in figure 2 where path [215] traverses five switching nodes before reaching the observation point [260] enabling this latter to receive a copy of all incoming traffic passing through port [231] and originated from Token Ring [230]. Set up of this path is done under the control of the network manager through the control point of each ATM switch with standard methods and techniques known from the art. What is different is that the control point of the switching node is instructed, by the overall network manager or whatever combination of tools and people in charge of keeping network up and running, to mark all cells arriving to port that must be snooped. Techniques for marking the cells are further discussed in the following. Whatever method is used all cells (irrespective of the channel identifier they are carrying) of incoming flow [418] from line [417], that is chosen to be snooped, are then marked [419] in blade [410]. Then, cells of flow [418] may have to be duplicated in the switch fabric itself (using the same overall mechanism as used for multicasting so that they are sent to the observation point through another blade [430] thus, creating an extra flow of cells [438], image of the snooped flow [418], and following the path reserved in the network for it up to the last blade

where measuring and observation equipment is connected. This duplication of the snooped flow may have also to take place in the transmit side of a blade whenever snooped flow and regular connections use the same blade for exiting a particular switching node. This case is illustrated here with blade [420]. Because of the multicast function that must be supported at switching nodes, already discussed in figure 3, cells belonging to some connections listed in table [426] may have also to be replicated locally multiple times e.g., on lines [421] and [424]. Therefore, snooped flow [428] must be replicated on top of the replication required by the multicasting of cells belonging to regular connections, if any is required. Consequently, this results in the sending of replicated marked cells possibly over the same line as the one borrowed by a regular connection. This is illustrated in this figure where snooped flow goes through line [421] too.

The here above description which focuses on the generation and handling of the marked cells needs only to be carried out in the source node where a port is snooped. All the other nodes, on the path to the observation point, automatically recognize and handle the marked cells to their final destination.

Figure 5 discusses how the cells can be marked so that they are recognized as snooped traffic and automatically routed, unaltered, to their final destination. ATM cell header [500] is 5-byte long including a one-byte CRC (called HEC) [510] to protect it thus, header carries four bytes i.e., 32 bits of information mainly comprised of a 16-bit VCI (Virtual Channel identifier) [520] and a 8-bit or 12-bit VPI (Virtual Path identifier) [530] depending upon the interface is said to be UNI (User to Network Interface) [540] or NNI (Network to Network Interface) [550]. Therefore, at least 24 bits are available within the header to identify a connection. Because, unlike of other protocols, this is not a universal address or identifier, it has only a local significance between two

switches and can freely be assigned by the network manager, there is no fear of a shortage of identifiers to differentiate the virtual connections even on the highest speed lines. The only constraint being that identifiers must be unique on the same physical communications line. Then, the simplest way of marking the cells is to borrow one bit of the VPi or VCi fields which, when turned on, means that cell belongs to the snooped traffic and must be processed accordingly. Alternatively, on the UNI interface, which has also a 4-bit GFC (Generic Flow Control) field [560] one of the bits could be devoted to marking the cells. Thus, in a preferred embodiment of the invention MSB (Most Significant Bit i.e., bit 15) of the VCI field [570] is chosen to mark the snooped cells. Hence, it becomes network manager responsibility not to use channel identifiers using this bit for the regular connections.

Figure 6 describes how cells are processed on the receive side of every blade. For each incoming cell [600] header is tested [605] so that to determine if cell has already been marked in a previous node at the origin of the snooped flow. If answer is positive then, next step [610] is to check if the current blade is indeed on the path of the snooped flow as defined by the network manager. If the answer is negative an error situation is encountered and cell must be discarded [611]. Normally, if snoop path has been correctly defined the answer to question [610] is positive in which case the marked cell can be forwarded directly to the switch fabric [630] after the necessary information has been added to it, so as cell can be steered through it and reach the output port from which it will be transported to a next node along the snoop path. What is appended to the cell to cross the switch fabric is highly dependent on the switch fabric design. However, it is generally broadly referred to as a 'Switch Tag' [615]. This latter comprises the necessary information to direct the incoming cell to the right port i.e., the outgoing snoop port [620] in this case. Thus, the processing of marked cells, arriving at

intermediate nodes, is expedited and those cells are moved along the snoop path unaltered until they reach the observation point.

However, if incoming cell [600] is not marked so that the
5 answer to question [605] is negative, it is a regular cell
(non-snoop cell) possibly belonging to one of the virtual
connection recognized on the receive side of the current
blade. Through the interrogation of a connection table [316],
shown in figure 3, it is determined if cell is actually
10 expected because it corresponds to a predefined connection. If
answer to question [625] is positive (the normal case) then,
cell follows the regular process i.e., a switch tag is added
[615] which specifies the outgoing connection port [635] so
that cell eventually traverses the switch fabric and reaches
15 the output port following connection it belongs to. However,
if current blade has been set, by the network manager, as
being the source of the snoop flow, snoop traffic must be
originated from this blade. Therefore, if answer to the next
question [640] is positive switch tag must be altered so that
20 switch fabric is instructed to replicate the regular cell to
the snoop port [645] too, eventually resulting in the forward-
ing of a marked cell. If current blade is not a snoop source
the answer to question [640] is negative and the above branch
is just bypassed. At this stage it is important to remark that
25 snoop port and connection port may, for some of the incoming
cells, match. In which case the replication of the regular
cell is not handled by the switch fabric itself but rather by
the transmit side of the outgoing blade. This is further
discussed in figure 7.

30 Finally, if answer to question [625] is negative, cell is
not recognized and should normally be discarded [650]. However,
if blade has been declared as the source of the snoop flow, so
that answer to question [655] is positive, all cells arriving
at the snooped port entry must be forwarded to the observation
35 point requiring that a switch tag be added [615] specifying the

snoop port [620] like previously described for the incoming marked cells.

Figure 7 focuses on the transmit side of a blade. Each regular cell received from a switch fabric port triggers the
5 interrogation of a local connection table [326] shown in figure 3. The objective of this interrogation is many fold. First, it must be decided to which output virtual connection cell should be sent. As already mentioned earlier this may include several lines, possibly all lines, implemented on the blade if cell
10 must be multicast. Therefore, a multicast mechanism allowing to replicate cells must be supported in the transmit side of every blade. Also, a VCi/VPi swap, a standard operation in ATM switches, is going to be performed. Since, in ATM networks, channel identifiers have only a local significance VPi and VCi
15 fields of cell headers are possibly exchanged at each node even though cells are following predetermined virtual paths. There are numerous ways, known from the art, of achieving these functions depending on the actual implementation of switches and blades. Irrespective of any particular implemen-
20 tation a standard method consists in adding, to the switch tag appended to the cells received from the lines in the receive side of each blade, a correlator i.e., an index used to address the local connection table found in each blade transmit side and from where all the information is found on what
25 to do with the current incoming cell. Namely, on what output port(s) current cell is to be sent and what VPi/VCi is to be used on each of them. This is generally carried out under the form of a linked list of destinations which is gone through each time a new cell is received so that it is replicated
30 accordingly. If no multicast is required i.e., cell is sent on a single outgoing port then, it just becomes a one-item list. Snooping per the invention takes advantage of the multicasting and is merged in this process whenever necessary as explained hereafter.

The algorithm for processing a cell, received from a switch fabric port, on the transmit side of every blade, first consist in testing a bit in the switch tag that was appended to the cell to know if cell must be 'recirculated' i.e., must
5 be replicated so as it participates to the generation of the snoop flow. It is worth noting here that switch fabric are not, per se, ATM switches but rather cell-switches. If switching nodes are indeed eventually performing ATM cell switching what it actually switched, in a switch fabric, is a true ATM
10 cell to which is appended enough information to handle it within the switch fabric itself and the blades. This is broadly refer to as switch tag in this description and was already discussed in previous figures. Switch tag carries information partly set by the receive side for the purpose of
15 multicasting cell on more than one switch fabric port. Also, it includes a correlator here above mentioned, used as an entry point into the local connection table, and more. Thus, among the overhead bits added to the ATM cell which is to be switched one bit is reserved for permitting a snoop recirculation of cells by the transmit side of the blades when necessary.
20 The first step is then to test if this bit [700] is set or not. When cell is received from the switch fabric this bit cannot be active since it is not set by the receive side or the switch fabric as it will be seen later. Thus, answer is
25 negative and one proceeds to the interrogation of the local connection table [710] to understand if cell is indeed expected and part of a predetermined flow. If answer is negative the next step [720] is to wonder if cell is marked, so as it belongs to the snoop flow. If answer is negative an error case
30 is detected, a logging of it is done and cell is discarded [730]. However, if answer is positive, a snoop cell (marked in current node or a previous node) is recognized and forwarded to the snoop port. This branch of the algorithm corresponds to what was depicted in figure 4 and flow [438]. If answer to
35 question [710] was however, positive in which case a cell belonging to a registered connection is to be processed the

next step [750] is the VPi/VCi swap, a standard operation in ATM network, previously discussed. After which a next question [760] is to determine if the current cell must be replicated. One way of achieving this, as already previously suggested, is to organize in the local connection table linked list of connections on which cells have to be replicated. Hence, unicast and multicast are handled alike i.e., when the end of the linked list is reached (which may comprise only one item) replication is stopped. Then, if one assumes first that answer to step [760] is negative a multicast recirculation bit is set [770] in the overhead bits of the cell (i.e., the switch tag) after which cell is forwarded [780] through the selected output port. Obviously, when cell is eventually sent over a telecommunications line, to which output port is tied, so as it is transported to another node, all the overhead (the switch tag) that was appended to the ATM cell for allowing internal processing, is removed. At this stage algorithm resumes in [700]. When, at first loop or during a subsequent loop (if cell is multicast), the end of the linked list is eventually reached and the answer to question [760] is positive then, the next step [790] is to wonder if a snoop cell must be generated too. If this is indeed the case the snoop recirculation bit is set [795]. Hence, when algorithm resume in [700] answer is finally positive so as branch ending in step [740] is gone through resulting in the forwarding of a snoop cell.

This Page Blank (uspto)

Claims:

1. A method of enabling remote surveillance of any entry port [231] to any switching node [202] of an ATM cell-relay network [220], said ATM cell-relay network comprising one or more of
5 said switching nodes each comprising a plurality of ports [110], said method comprising the steps of:
- setting up a path [215], through said ATM network, between said entry port [231] and a remote observation point [260];
- duplicating [428] [438] all cells of the incoming traffic
10 entering through said entry port [417], said step of duplicating cells further including the step of:
- marking [419] all said duplicated cells;
- transporting, along said path, said marked duplicated cells up to said observation point.
- 15 2. The method according to claim 1 wherein the step of marking all duplicated cells further includes the steps of:
- reserving one bit [570] of the cells channel identifier to unambiguously distinguish said duplicated cells;
- asserting said reserved bit of every said duplicated cell.
- 20 3. The method according to claim 2 wherein the step of reserving one bit alternatively utilizes one bit of the cells GFC (Generic Flow Control) field [560].
4. The method according to claim 2 wherein the step of reserving one bit utilizes the most significant bit of VCi (Virtual
25 Channel identifier) field part of the cells channel identifier field.

5. The method according to any one of the previous claims wherein said duplicated cells are marked with any unique combination of bits within header [500] of said cells.

6. The method according to any one of the previous claims
5 wherein the step of duplicating cells is carried out within said switching node [202] of said entry port [231].

7. The method according to any one of the previous claims wherein the step of duplicating cells includes duplicating:

unexpected cells [625];
10 errored cells;
nonconforming cells.

8. The method according to any one of the previous claims wherein the step of transporting said marked duplicated cells includes, in all intermediate said switching nodes along said
15 path up to said observation point, the further steps of:

recognizing immediately [605] said marked duplicated cells when entering each intermediate said switching node;
testing [610] if said port of each intermediate said switching node is indeed on said path to said observation path;

20 if yes:

skipping the regular ATM cell processing;
keeping unaltered said marked duplicated cells;
moving forward [620] said marked duplicated cell;

if no:

25 discarding [611] said marked duplicated cells;
recording an error.

9. A system, in particular an ATM cell-relay network surveillance system, comprising means adapted for carrying out the method according to any one of the previous claims.

10. The system according to claim 9 further comprising, in each
5 of said switching node:

a cell switch fabric [300];

adaptive blades [310] between said ports and said cell
switch fabric.

11. The system according to claim 10 further comprising, in
10 said cell switch fabric and/or in said adaptive blades means
for replicating cells.

This Page Blank (uspto)

**SYSTEM AND METHOD FOR ENABLING
REMOTE SURVEILLANCE
OF ATM NETWORK SWITCHING NODE PORTS**

Abstract

5 In an ATM cell-relay network usually comprising several
switching nodes a method enabling remote surveillance of any
entry port to any switching node of the network is disclosed.
The method first assumes that a path is set up from the entry
port to a remote observation point. Then, all cells of the
10 incoming traffic, entering entry port to be watched, are
duplicated and marked. After which they are transported,
unaltered, following the path, up to the observation point.
The invention permits that any entry port of an ATM network,
that may well span over large geographic areas, be conven-
15 iently observable and analyzed transparently i.e., without
disturbing users traffic, from a remote location so as network
can be maintained and run trouble free.

Figure 2.

20

This Page Blank (uspto)

FR 9 99 080
AZNAR et al.
1/7




Figure 1

FR 9 99 080

AZNAR et al.

2/7

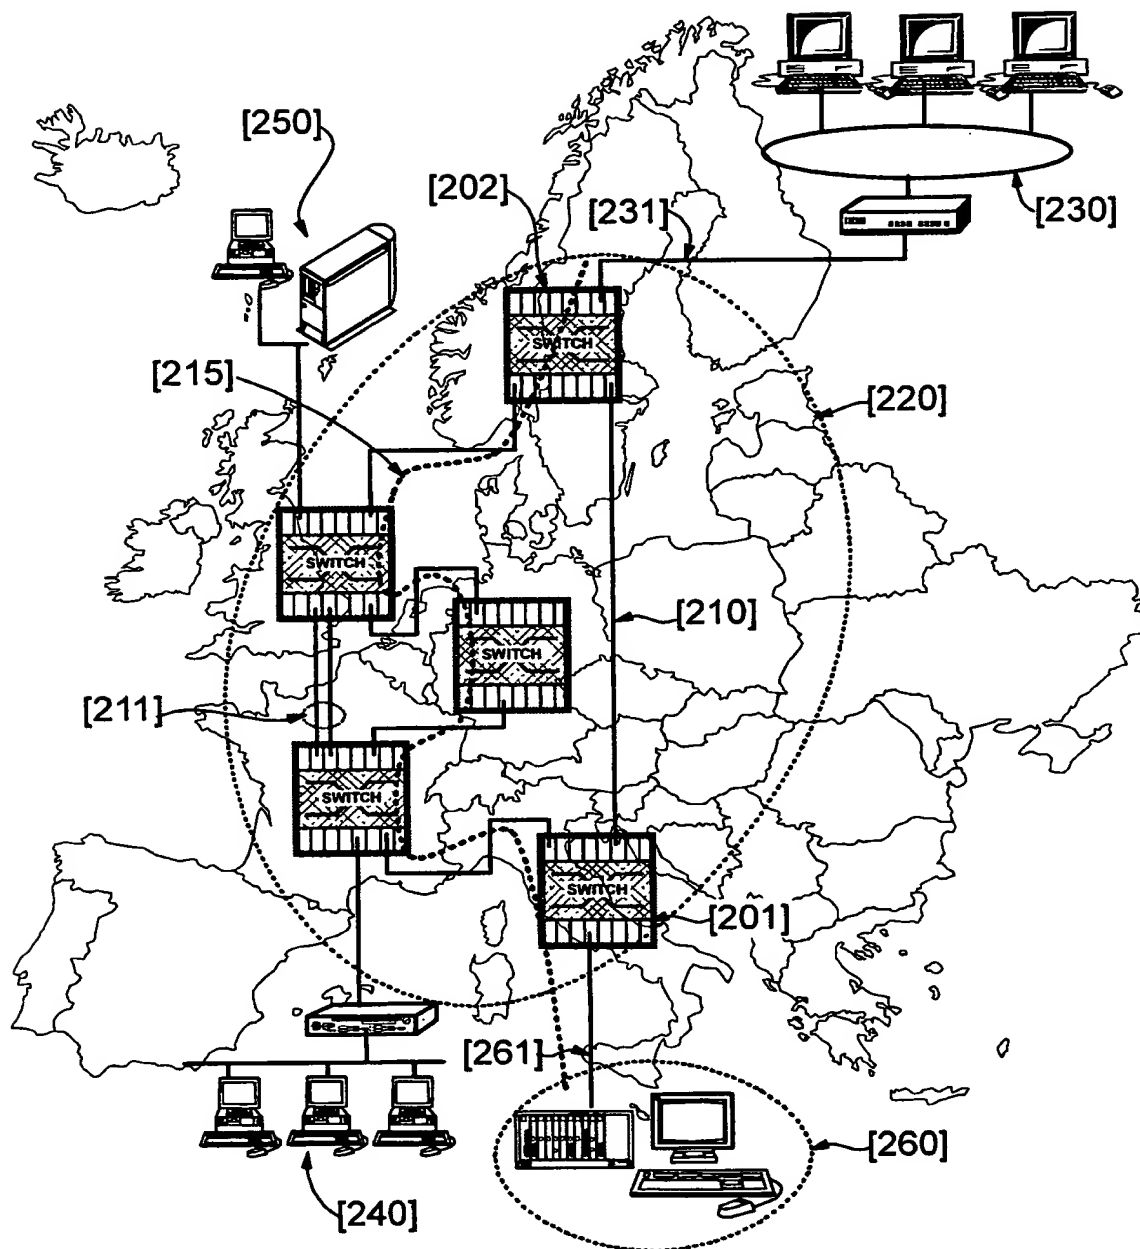


Figure 2

FR 9 99 080
AZNAR et al.
3/7




Figure 3

FR 9 99 080
AZNAR et al.
4/7




Figure 4

FR 9 99 080
AZNAR et al.
5/7




Figure 5

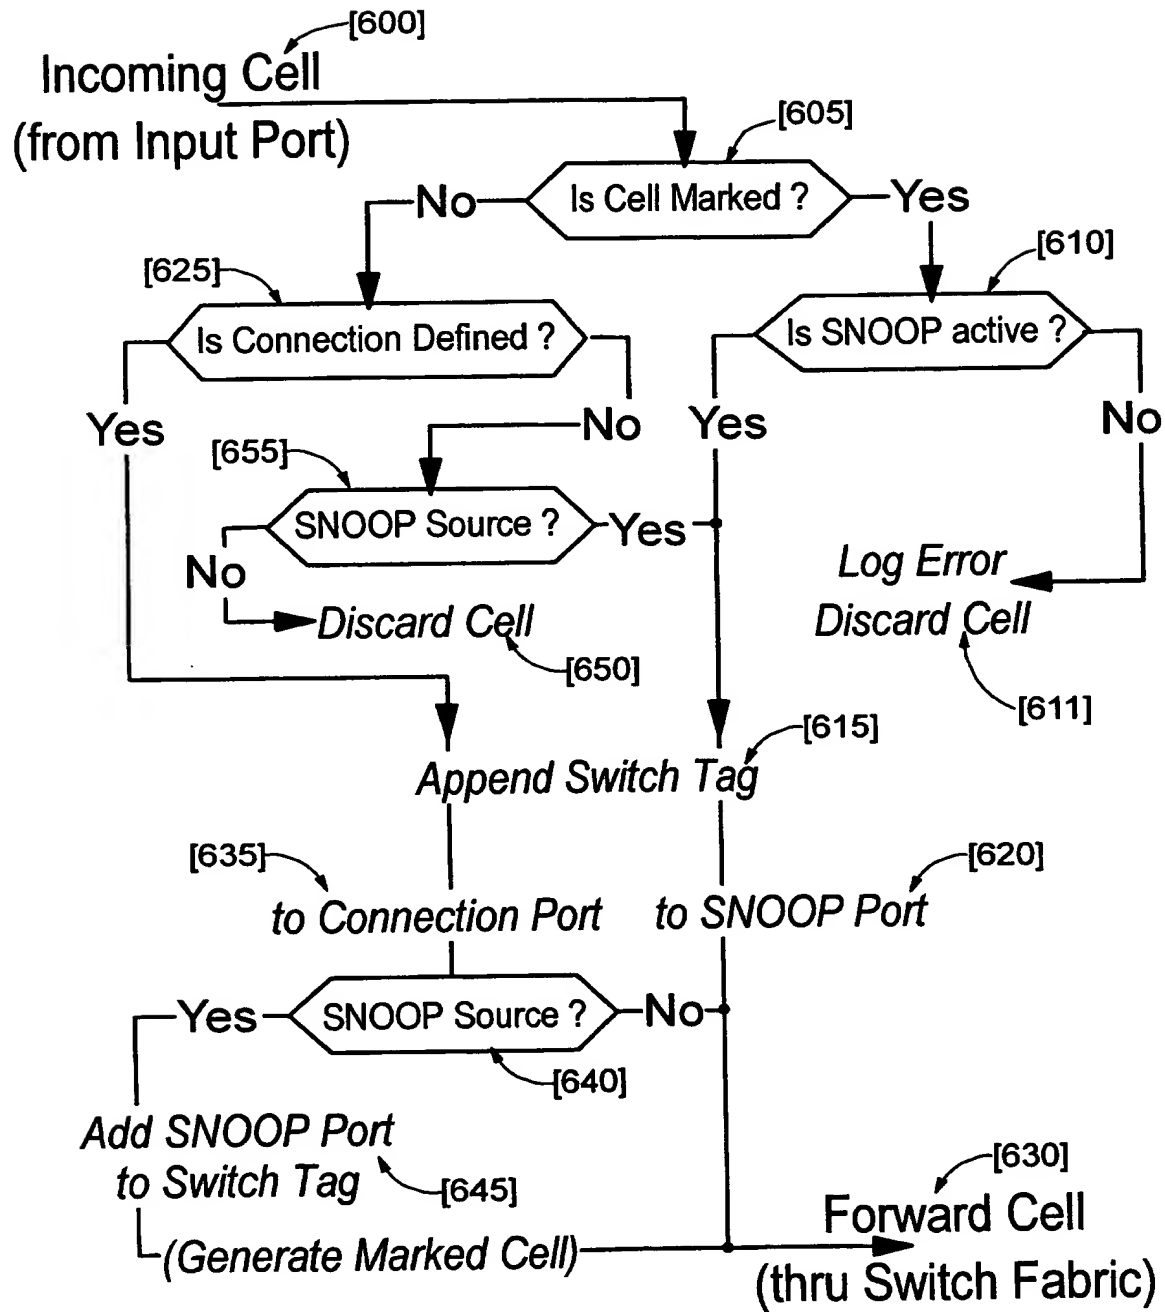


Figure 6

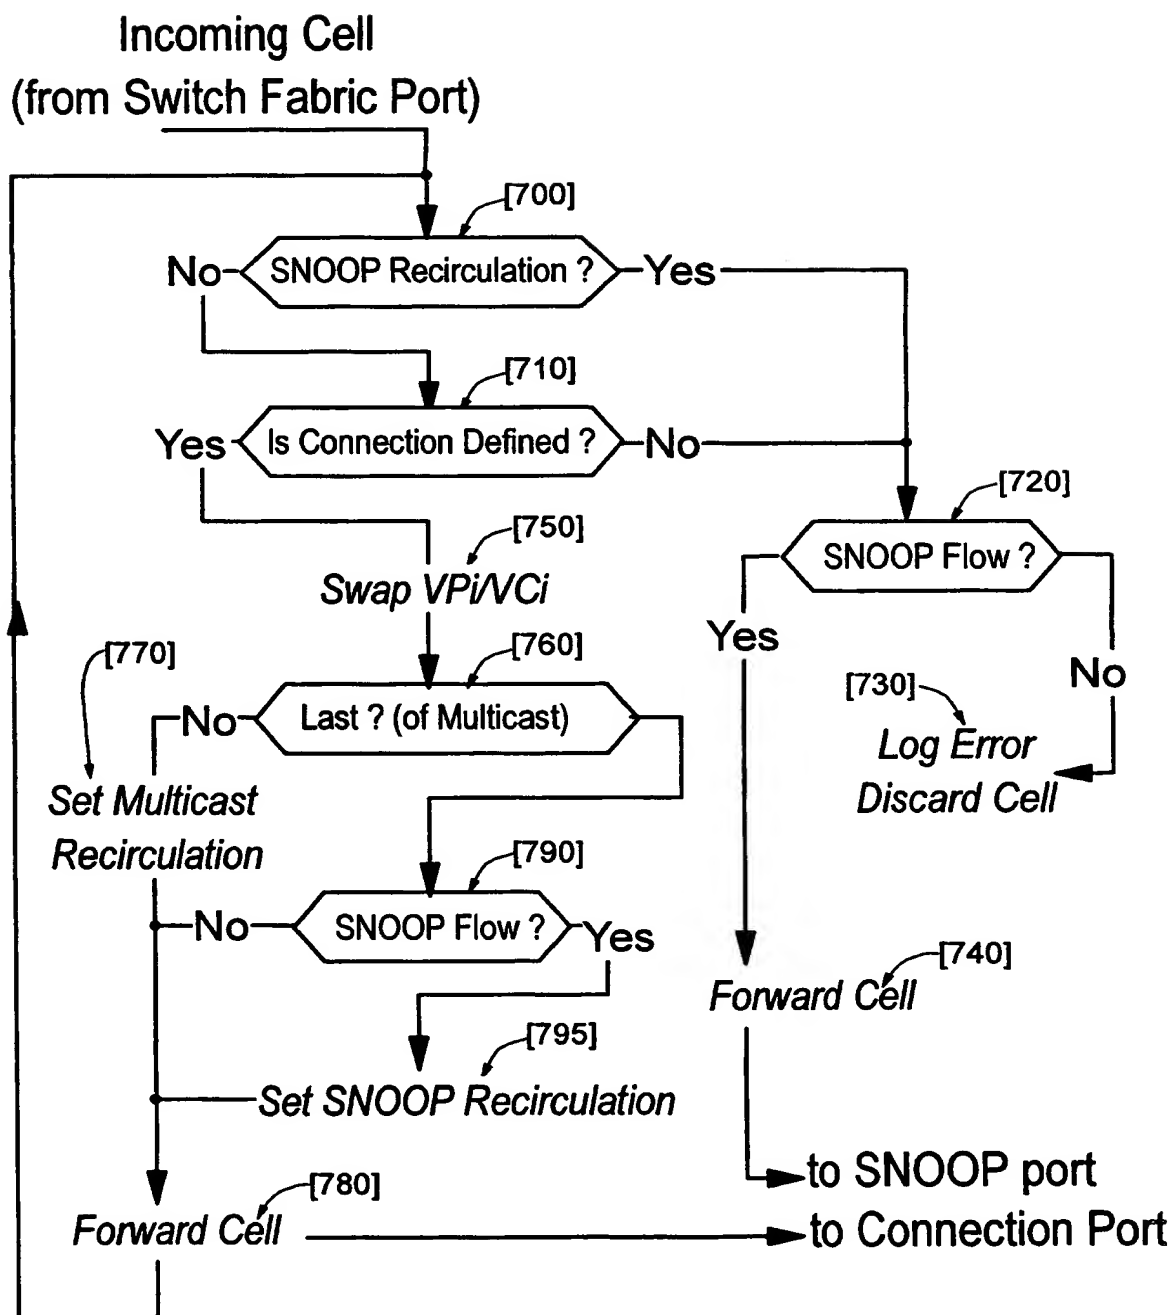


Figure 7

This Page Blank (uspto)